arXiv:1601.02068v4 [stat.ML] 2 Dec 2016

On Computationally Tractable Selection of Experiments in Regression Models

Yining Wang[1], Adams Wei Yu[1], and Aarti Singh[1]

[1]Machine Learning Department, Carnegie Mellon University

December 5, 2016

Abstract

We derive computationally tractable methods to select a small subset of experiment settings from a large pool of given design points. The primary focus is on linear regression models, while the technique extends to generalized linear models and Delta's method (estimating functions of linear regression models) as well. The algorithms are based on a continuous relaxation of an otherwise intractable combinatorial optimization problem, with sampling or greedy procedures as post-processing steps. Formal approximation guarantees are established for both algorithms, and simulations on synthetic data confirm the effectiveness of the proposed methods.

Keywords: optimal selection of experiments, A-optimality, computationally tractable methods

1 Introduction

Consider the linear regression model

$$y = X\beta_0 + \varepsilon,$$

where $X \in \mathbb{R}^{n\times p}$ is the design matrix, $y \in \mathbb{R}^n$ is the response and $\varepsilon \sim \mathcal{N}_n(0, \sigma^2 I_n)$ are homoscedastic Gaussian noise with variance σ^2. β_0 is a p-dimensional regression model that one wishes to estimate. We consider the "large-scale low-dimensional" setting where both $n, p \to \infty$ and $p < n$, and X has full column rank. A common estimator is the *Ordinary Least Squares* (OLS) estimator:

$$\hat{\beta}^{\text{ols}} = \text{argmin}_{\beta\in\mathbb{R}^p} \|y - X\beta\|_2^2 = (X^\top X)^{-1}X^\top y.$$

Despite the simplicity and wide applicability of OLS, in practice it may not be desirable to obtain the full n-dimensional response vector y due to measurement constraints. It is then a common practice to select a small subset of rows (e.g., $k \ll n$ rows) in X so that the statistical efficiency of regression on the selected subset of design points is maximized. Compared to the classical *experimental design* problem (Pukelsheim, 1993) where X can be freely designed, in this work we consider the setting where the selected design points must come from an existing (finite) design pool X. Below we list three example applications where such constraints are relevant:

Example 1 (Wind speed prediction) In Chen et al. (2015) a data set is created to record wind speed across a year at established measurement locations on highways in Minnesota, US. Due to instrumentation constraints, wind speed can only be measured at intersections of high ways, and a small subset of such intersections is selected for wind speed measurement in order to reduce data gathering costs.

Example 2 (CPU benchmarking) Central processing units (CPU) are vital to the performance of a computer system. It is an interesting statistical question to understand how known manufacturing parameters (clock period, cache size, etc.) of a CPU affect its execution time (performance) on benchmark computing tasks (Ein-Dor & Feldmesser, 1987). As the evaluation of real benchmark execution time is time-consuming and costly, it is desirable to select a subset of CPUs available in the market with diverse range of manufacturing parameters so that the statistical efficiency is maximized by benchmarking for the selected CPUs.

Example 3 (Neural visual perception) Understanding how humans perceive visual objects is an important question in psychological and neural science. In an experimental study (Leeds et al., 2014) a group of human subjects are shown several pictures of visual objects such as cars, toys, dogs and neural-biological responses are recorded. As experimental studies are expensive, it is an important question to select a subset of "representative" visual stimuli in order to reduce experimental cost. Again, visual stimuli cannot be arbitrarily designed and must be selected from an existing pool of natural pictures.

In this work, we focus on **computationally tractable** methods for experiment selection that achieves near-optimal statistical efficiency in linear regression models. We consider two experiment selection models: the *with replacement* model where each design point (row of X) can be selected more than once with independent noise involved for each selection, and the *without replacement* model where distinct row subsets are required. We propose two computationally tractable algorithms: one sampling based algorithm that achieves $1 + o(1)$ approximation of the statistically optimal solution for the with replacement model and $O(\log k)$ approximation for the without replacement model, assuming that the design pool X is well-conditioned. In the case of ill-conditioned design, we pose a greedy method that achieves $O(p^2/k)$ approximation for the without replacement model.

2 Problem formulation and backgrounds

We first give a formal definition of the experiment selection problem in linear regression models:

Definition 2.1 (experiment selection models)**.** *Let X be a known $n \times p$ design matrix with full column rank and k be the subset budget, with $p \le k \le n$. An experiment selection method find a subset $S \subseteq [n]$ of size k and then $y_S = X_S\beta_0 + \tilde{\varepsilon}$ is observed, where each coordinate of $\tilde{\varepsilon}$ is i.i.d. Gaussian random variable with zero mean and equal covariance. Two scenarios are considered:*

1. With replacement*: S is a multi-set which allows duplicates of indices. Note that fresh (independent) noise is imposed after experiment selection, and hence duplicate design points have independent noise.*

Table 1: Summary of approximation results. $\Sigma_* = X^\top \mathrm{diag}(\pi^*) X$, where π^* is the optimal solution of Eq. (3).

Algorithm	Model	$C(n,p,k)$	Assumptions
sampling	with rep.	$1+o(1)$	$\|\Sigma_*^{-1}\|_2 \|X\|_\infty^2 \log k \to 0$
sampling	without rep.	$O(\log k)$	$\|\Sigma_*^{-1}\|_2 \|X\|_\infty^2 \log k \to 0$
greedy	without rep.	**Rigorous**: $1+\frac{p(p+1)}{2(k-p+1)}$ **Conjecture**: $1+O(\frac{p}{k-p})$	$k > p$

2. Without replacement*: S is a standard set that may not have duplicate indices.*

As evaluation criterion, we consider the mean square error $\mathbb{E}\|\hat{\beta}_k - \beta_0\|_2^2$, where $\hat{\beta}_k = (X_S^\top X_S)^{-1} X_S^\top y_S$ is the OLS estimate on the selected design X_S and its corresponding response y_S. In the standard setting of $n > p$ and $\mathrm{rank}(X) = p$, it is well-known that the statistically optimal subset S^* is given by the *A-optimality* criterion (Pukelsheim, 1993):

$$S^* = \mathrm{argmin}_{S \subseteq [n], |S| \le k} \mathrm{tr}\left[\left(X_S^\top X_S\right)^{-1}\right]. \tag{1}$$

Here $S^* \subseteq [n]$ is a multi-set if with-replacement model is considered and a standard set if without-replacement model is used.

Despite the statistical optimality of Eq. (1), the optimization problem is combinatorial in nature and the optimal subset S^* is difficult to compute. A brute-force search over all possible subsets of size k requires $O(n^k k^3)$ operations, which is infeasible for even moderately sized designs X.

In this paper we consider computationally-feasible methods that approximately compute the optimal design in Eq. (1). By "computationally feasible" or "computationally tractable", we refer to algorithms whose running time scales *polynomially* with input dimension n, p and k. For example, the OLS estimate $\hat{\beta}_k = (X_S^\top X_S)^{-1} X_S^\top y_S$ takes $O(kp^2 + p^3)$ operations to compute and is computationally feasible. In contrast, the brute-force search method for optimizing Eq. (1) requires $O(n^k k^3)$ running time, which is exponential in k when both p, k increases with n and thus computationally intractable. Let $\hat{S}$ be the subset produced by such a computationally efficient algorithm. Due to computational constraints, $\hat{S}$ may not achieve the optimal statistical efficiency as S^* defined in Eq. (1). Nevertheless, we wish to guarantee the statistical property of $\hat{S}$ in the following sense:

$$F(\hat{S}; X) \le C(n,p,k) \cdot F(S^*; X) \quad \text{where} \quad F(S; X) = \mathrm{tr}\left[\left(X_S^\top X_S\right)^{-1}\right]. \tag{2}$$

Here $C(n,p,k) \ge 1$ is the *approximation ratio* of $\hat{S}$. In Table 1 we summarize the approximation ratio of computationally-tractable experiment selection algorithms presented in this paper.

2.1 Related work

There has been an increasing amount of work on fast solvers for the general least-square problem $\min_\beta \|y - X\beta\|_2^2$. Most of existing work along this direction (Woodruff, 2014; Dhillon et al., 2013; Drineas et al., 2011; Raskutti & Mahoney, 2015) focuses solely on the computational aspects and do

not consider statistical constraints such as limited measurements of y. A convex optimization formulation was proposed in Davenport et al. (2015) for a constrained adaptive sensing problem, without finite sample guarantees on the combinatorial problem. In Horel et al. (2014) a computationally tractable approximation algorithm was proposed for the D-optimality criterion of the experimental design problem. However, the core idea in Horel et al. (2014) of pipage rounding an SDP solution (Ageev & Sviridenko, 2004) is not appplicable to A-optimality type criterion as defined in Eq. (1).

Popular subsampling techniques such as leverage score sampling (Drineas et al., 2008) were studied in least square and linear regression problems (Zhu et al., 2015; Ma et al., 2015; Chen et al., 2015). While computation remains the primary subject, measurement constraints and statistical properties were also analyzed within the linear regression model (Zhu et al., 2015). However, few existing (computationally efficient) methods achieve optimal statistical efficiency in terms of estimating the underlying linear model β_0.

Another related area is *active learning* (Chaudhuri et al., 2015; Hazan & Karnin, 2015; Sabato & Munos, 2014), which is a stronger setting where feedback from prior measurements can be used to guide subsequent data selection. Chaudhuri et al. (2015) analyzes an SDP relaxation in the context of active maximum likelihood estimation. However, the analysis in Chaudhuri et al. (2015) only works for the with replacement model and the two-stage feedback-driven strategy proposed in Chaudhuri et al. (2015) is not available under the experiment selection model defined in Definition 2.1.

2.2 Notations

For a matrix $A \in \mathbb{R}^{n\times m}$, we use $\|A\|_p = \sup_{x\neq 0} \frac{\|Ax\|_p}{\|x\|_p}$ to denote the induced p-norm of A. In particular, $\|A\|_1 = \max_{1\leq j\leq m} \sum_{i=1}^n |A_{ij}|$ and $\|A\|_\infty = \max_{1\leq i\leq n} \sum_{j=1}^m |A_{ij}|$. $\|A\|_F = \sqrt{\sum_{i,j} A_{ij}^2}$ denotes the Frobenius norm of A. Let $\sigma_1(A) \geq \sigma_2(A) \geq \cdots \geq \sigma_{\min(n,m)}(A) \geq 0$ be the singular values of A, sorted in descending order. The condition number $\kappa_2(A)$ is defined as $\kappa(A) = \sigma_1(A)/\sigma_{\min(n,m)}(A)$. For sequences of random variables X_n and Y_n, we use $X_n \xrightarrow{p} Y_n$ to denote X_n converges in probability to Y_n. We say $a_n \lesssim b_n$ if $\lim_{n\to\infty} |a_n/b_n| \leq 1$ and $a_n \gtrsim b_n$ if $\lim_{n\to\infty} |b_n/a_n| \leq 1$.

3 Methods and main results

We describe two computationally feasible algorithms for the experiment selection problem, both based on a continuous relaxation of the combinatorial A-optimality problem. Statistical efficiency bounds are present for both algorithms, with detailed proofs given in Sec. 6.

3.1 Continuous relaxation of Eq. (1)

Consider the following continuous optimization problem:

$$\pi^* = \operatorname{argmin}_{\pi\in\mathbb{R}^p} f(\pi; X) = \operatorname{argmin}_{\pi\in\mathbb{R}^p} \operatorname{tr}\left[\left(X^\top \operatorname{diag}(\pi) X\right)^{-1}\right], \tag{3}$$
$$s.t. \quad \pi \geq 0, \quad \|\pi\|_1 \leq k,$$
$$\|\pi\|_\infty \leq 1 \quad \text{(only for the without replacement model)}.$$

Note that the $\|\pi\|_\infty \le 1$ constraint is only relevant for the without replacement model and for the with replacement model we drop this constraint in the optimization problem. It is easy to verify that both the objective funnction $f(\pi; X)$ and the feasible set in Eq. (3) are convex, and hence the global optimal solution π^* of Eq. (3) can be obtained using computationally tractable algorithms. In particular, we describe an SDP formulation and a practical projected gradient descent algorithm in Appendix B and the supplementary material, both provably converge to the global optimal solution of Eq. (1) with running time scaling polynomially in n, p and k.

The following Gaussian design example gives an intuitive interpretation of the optimization problem Eq. (3) and show in what way it improves over the trivial experiment selection method of choosing rows of X uniformly at random.

Example: Gaussian design We use a simple Gaussian example to show that non-uniform weights $\pi_i \neq k/n$ could improve the objective function $f(\pi; X)$. Suppose $x_1, \cdots, x_n \overset{i.i.d.}{\sim} \mathcal{N}_p(0, \Sigma_0)$. Let π^{unif} be the uniformly weighted solution of $\pi_i^{\text{unif}} = k/n$, corresponding to selecting each row of X uniformly at random. We then have

$$f(\pi^{\text{unif}}; X) = \frac{1}{k}\text{tr}\left[\left(\frac{1}{n}X^\top X\right)^{-1}\right] \xrightarrow{p} \frac{1}{k}\text{tr}(\Sigma_0^{-1}).$$

On the other hand, let $B^2 = \gamma \text{tr}(\Sigma_0)$ for some universal constant $\gamma > 1$ and $\mathbb{B} = \{x \in \mathbb{R}^p : \|x\|_2^2 \le B^2\}$, $\mathbb{X} = \{x_1, \cdots, x_n\}$. By Markov inequality, $|\mathbb{X} \cap \mathbb{B}| \gtrsim \frac{\gamma n}{1-\gamma}$. Define weighted solution π^w as $\pi_i^w \propto 1/p(x_i|\mathbb{B}) \cdot I[x_i \in \mathbb{B}]$ normalized such that $\|\pi_i^w\|_1 = k$. [1] Then

$$f(\pi^w; X) = \frac{1}{k}\frac{\text{tr}\left[\left(\frac{1}{n}\sum_{x_i \in \mathbb{X}\cap\mathbb{B}} \frac{x_i x_i^\top}{p(x_i|\mathbb{B})}\right)^{-1}\right]}{\frac{1}{n}\sum_{x_i \in \mathbb{X}\cap\mathbb{B}} 1/p(x_i|\mathbb{B})}$$

$$\xrightarrow{p} \frac{1}{k}\frac{n}{|\mathbb{X}\cap\mathbb{B}|}\frac{\text{tr}\left[\left(\int_{\mathbb{B}} xx^\top \mathrm{d}x\right)^{-1}\right]}{\int_{\mathbb{B}} 1 \mathrm{d}x} \lesssim \frac{p^2}{(1-\gamma)\text{tr}(\Sigma_0)}.$$

Here in the last inequality we apply Lemma A.1. Because $\frac{\text{tr}(\Sigma_0^{-1})}{p} = \frac{1}{p}\sum_{i=1}^p \frac{1}{\sigma_i(\Sigma_0)} \ge \left(\frac{1}{p}\sum_{i=1}^p \sigma_i(\Sigma_0)\right)^{-1} = \frac{p}{\text{tr}(\Sigma_0)}$ by Jensen's inequality, we conclude that in general $f(\pi^w; X) < f(\pi^{\text{unif}}; X)$, and the gap is larger for ill-conditioned covariance Σ_0. This example shows that uneven weights in π helps reducing the trace of inverse of the weighted covariance $X^\top \text{diag}(\pi) X$.

We also present several facts about the continuous optimization problem. Proofs are trivial and are given in Sec. 6.1.

Fact 3.1. $f(\pi^*; X) \le F(S^*; X)$.

Fact 3.2. *Let π and π' be feasible solutions of Eq. (3) such that $\pi_i \le \pi_i'$ for all $i = 1, \cdots, n$. Then $f(\pi; X) \ge f(\pi'; X)$, with equality if and only if $\pi = \pi'$*

[1] Note that π^* may *not* be the optimal solution of Eq. (3); however, it suffices for the purpose of the demonstration of improvedment resulting from non-uniform weights.

Fact 3.3. $\|\pi^*\|_1 = k$.

From Fact 3.1, it is clear the optimal value of Eq. (3) *lower* bounds the combinatorial optimal solution in Eq. (1). However, the n-dimensional vector π^* is *not* a valid experiment selection for X without further processing: π^* may have uneven weights and more than k non-zero coordinates. It is thus important to develop post-processing procedures that transform π^* into a valid subset solution $\hat{S}$ without losing much of the statistical efficiency in π^*. This is the topic of the next two sections.

3.2 Sampling based experiment selection

A natural idea of obtaining a valid subset S of size k is by *sampling* from a weighted row distribution specified by π^*. More specifically, let $\hat{S} = \{i_1, \cdots, i_k\}$ where $i_1, \cdots, i_k$ are sampled, either with or without replacement depending on the model, from the following distribution:

$$\Pr[i_t = j] = \frac{\pi^*_j}{k}, \qquad t = 1, \cdots, k; j = 1, \cdots, n.$$

The sampling procedure is easy to understand in an asymptotic sense. For example, consider the with replacement model where $i_1, \cdots, i_k$ are i.i.d. sampled from the distribution $\Pr[i_t = j] = \pi^*_j/k$. It is easy to verify that $\mathbb{E}X_{\hat{S}}^\top X_{\hat{S}} = X^\top \mathrm{diag}(\pi^*)X$. By weak law of large numbers, $X_{\hat{S}}^\top X_{\hat{S}} \xrightarrow{p} X^\top \mathrm{diag}(\pi^*)X$ as $k \to \infty$ and hence $F(\hat{S}; X) \xrightarrow{p} f(\pi^*; X)$ by continuous mapping theorem. A more refined analysis is presented in Theorem 3.1 to provide explicit conditions under which the asymptotic approximations are valid and on the statistical efficiency of $\hat{S}$ as well as analysis under the more restrictive without replacement regime.

Theorem 3.1. *Suppose* $\|\Sigma_*^{-1}\|_2 \|X\|_\infty^2 \log k \to 0$ *as* $n, k \to \infty$, *where* $\Sigma_* = X^\top \mathrm{diag}(\pi^*)X$. *Then with probability* $1 - o(1)$ *the subset* $\hat{S}$ *satisfies* $F(\hat{S}; X) \le (1 + o(1))F(S^*; X)$ *for the with replacement model and* $F(\hat{S}; X) \le O(\log k) \cdot F(S^*; X)$ *for the without replacement model.*

Though it is tempting to establish the proximity betweeen $X_{\hat{S}}^\top X_{\hat{S}}$ and Σ_* directly using matrix concentration results, such an approach leads to undesired dependency over $\kappa(\Sigma_*)$ in the conditions of n and k. In order to prove theorem 3.1 under weaker conditions, we adapt the proof technique of Spielman and Srivastava in their seminal work on spectral sparsification of graphs (Spielman & Srivastava, 2011). More specifically, we prove the following stronger result which shows that $X_{\hat{S}}^\top X_{\hat{S}}$ is a *spectral approximation* of Σ_* with high probability, under suitable conditions.

Lemma 3.1. *Let* $\hat{S}$ *be constructed by sampling* with *replacement and define* $\Sigma_{\hat{S}} = X_{\hat{S}}^\top X_{\hat{S}}$. *Suppose* $\|\Sigma_*^{-1}\|_2 \|X\|_\infty^2 \log k \to 0$. *Then with probability* $1 - o(1)$ *the following holds* uniformly *over all* $z \in \mathbb{R}^p$*:*

$$z^\top \Sigma_{\hat{S}} z = (1 + o(1)) z^\top \Sigma_* z.$$

Lemma 3.1 implies that $\sigma_j(\Sigma_{\hat{S}}) = (1+o(1))\sigma_j(\Sigma_*)$ for all $j = 1, \cdots, p$. Recall that $F(\hat{S}; X) = \sum_{j=1}^p \sigma_j(\Sigma_{\hat{S}})^{-1}$ and $f(\pi^*; X) = \sum_{j=1}^p \sigma_j(\Sigma_*)$. The with replacement result of Theorem 3.1 is then a simple consequence of Lemma 3.1 because $F(\hat{S}; X) \le (1 + o(1))f(\pi^*; X) \le (1 + o(1))F(S^*; X)$.

To establish the without replacement result in Theorem 3.1, we prove the following lemma that exploits the additional infinity norm constraint in Eq. (3):

Lemma 3.2. *Let $\hat{S}$ be constructed by sampling* with *replacement and suppose* $\|\pi^*\|_\infty \le 1$. *Then* $\max_{1\le i\le n} \#\{i \in \hat{S}\} \le O_P(\log k)$.

Let $\tilde{S} \subseteq \hat{S}$ be the set of distinct elements in $\hat{S}$. Lemma 3.2 shows that, with probability $1-o(1)$, $|\tilde{S}| \ge |\hat{S}|/O(\log k)$, which further implies that $F(\tilde{S}; X) \le O(\log k) \cdot F(\hat{S}; X)$. Complete proofs of Lemma 3.1, 3.2 and Theorem 3.1 are placed in Sec. 6.

To provide further insights into the condition $\|\Sigma_*^{-1}\|_2 \|X\|_\infty^2 \log k \to 0$ in Theorem 3.1, we consider a Gaussian design example and derive interpretable conditions.

Example (Gaussian design) Suppose $x_1, \cdots, x_n \overset{i.i.d.}{\sim} \mathcal{N}_p(0, \Sigma_0)$. Then $\|X\|_\infty^2 \le O_P(\|\Sigma_0\|_2^2 p \log n)$. In addition, by simple algebra $\|\Sigma_*^{-1}\|_2 \le \min\{1, p^{-1}\kappa(\Sigma_*)\}\mathrm{tr}(\Sigma_*^{-1}) \le \min\{1, p^{-1}\kappa(\Sigma_*)\}F(S^*; X)$. Using a very conservative upper bound of $F(S^*; X)$ by sampling rows in X uniformly at random and apply weak law of large numbers and the continuous mapping theorem, we have that $F(S^*; X) \lesssim \frac{1}{k}\mathrm{tr}(\Sigma_0^{-1})$. In addition, $\mathrm{tr}(\Sigma_0^{-1})\|\Sigma_0\|_2 \le p\|\Sigma_0^{-1}\|_2\|\Sigma_0\|_2 = p\kappa(\Sigma_0)$. Subsequently, the condition $\|\Sigma_*^{-1}\|_2\|X\|_\infty^2 \log k \to 0$ is implied by

$$\frac{p\kappa(\Sigma_*)\kappa(\Sigma_0)\log n \log k}{k} \to 0. \tag{4}$$

Essentially, the condition is reduced to $k \gtrsim \kappa(\Sigma_0)\kappa(\Sigma_*) \cdot p \log n \log k$. The linear dependency on p is necessary, as we consider the low-dimensional linear regression problem and $k < p$ would imply an infinite mean-square error in estimation of β_0. We also remark that the condition is scale-invariant, as $X' = \xi X$ and X share the same quantity $\|\Sigma_*^{-1}\|_2\|X\|_\infty^2 \log k$.

3.3 Greedy experiment selection

In Avron & Boutsidis (2013) an interesting upper bound on $F(S^*; X)$ for the without replacement model (i.e., S^* does not have duplicate indices) is proved, which we summarize below:

Lemma 3.3 (Avron & Boutsidis (2013), Lemma 3.9, simplified)**.** $F(S^*; X) \le \frac{n-p+1}{k-p+1}\mathrm{tr}[(X^\top X)^{-1}]$ *for $p < k \le n$, under the without replacement model.*

One consequence of Lemma 3.3 is a computationally tractable greedy removal procedure, which is summarized below:

input : $X \in \mathbb{R}^{n\times p}$, Initial subset $S_0 \subseteq [n]$, target size $k \le |S_0|$.
output: $\hat{S} \subseteq [n]$, a selected subset of size k.
Initialization: $t = 0$.
1. Find $j^* \in S_t$ such that $\mathrm{tr}[(X_{S_t\setminus\{j^*\}}^\top X_{S_t\setminus\{j^*\}})^{-1}]$ is minimized.
2. Remove j^* from S_t: $S_{t+1} = S_t\setminus\{j^*\}$.
3. Repeat steps 1 and 2 until $|S_t| = k$. Output $\hat{S} = S_t$.

Figure 1: Greedy based experiment selection.

The following corollary immediately follows Lemma 3.3:

Corollary 3.1. $F(\hat{S}; X) \le \frac{|S_0|-p+1}{k-p+1}F(S_0; X)$, *for $p < k \le |S_0|$.*

In Avron & Boutsidis (2013) the greedy removal procedure in Figure 1 is applied to the entire design set $S_0 = [n]$, which gives approximation guarantee $F(\hat{S}; X) \leq \frac{n-p+1}{k-p+1}\mathrm{tr}[(X^\top X)^{-1}]$. This results in an approximation ratio of $\frac{n-p+1}{k-p+1}$ as defined in Eq. (2) by applying the trivial bound $\mathrm{tr}[(X^\top X)^{-1}] \leq F(S^*; X)$, which is tight for a design that has exactly k non-zero rows.

To further improve the approximation ratio, we consider applying the greedy removal procedure with S_0 equal to the support of π^*; that is, $S_0 = \{j \in [n] : \pi^*_j > 0\}$. Because $\|\pi^*\|_\infty \leq 1$ under the without replacement setting, we have the following corollary:

Corollary 3.2. *Let S_0 be the support of π^* and suppose $\|\pi^*\|_\infty \leq 1$. Then $F(\hat{S}; X) \leq \frac{\|\pi^*\|_0-p+1}{k-p+1} f(\pi^*; X) \leq \frac{\|\pi_*\|_0-p+1}{k-p+1} F(S^*; X)$.*

It is thus important to upper bound the support size $\|\pi^*\|_0$. With the trivial bound of $\|\pi^*\|_0 = n$ we recover the $\frac{n-p+1}{k-p+1}$ approximation ratio by applying Figure 1 to $S_0 = [n]$. In order to bound $\|\pi^*\|_0$ away from n, we consider the following assumption imposed on X:

Assumption 3.1. *Define mapping $\phi : \mathbb{R}^p \to \mathbb{R}^{\frac{p(p+1)}{2}}$ as $\phi(x) = (\xi_{ij} x(i)x(j))_{1\leq i\leq j\leq p}$, where $x(i)$ denotes the ith coordinate of a p-dimensional vector x and $\xi_{ij} = 1$ if $i = j$ and $\xi_{ij} = 2$ otherwise. Denote $\tilde{\phi}(x) = (\phi(x), 1) \in \mathbb{R}^{\frac{p(p+1)}{2}+1}$ as the affine version of $\phi(x)$. For any $\frac{p(p+1)}{2} + 1$ distinct rows of X, their mappings under $\tilde{\phi}$ are linear independent.*

Assumption 3.1 is essentially a general-position assumption, which assumes that no $\frac{p(p+1)}{2} + 1$ design points in X lie on a degenerate affine subspace after a specific quadratic mapping. Like other similar assumptions in the literature (Tibshirani, 2013), Assumption 3.1 is very mild and almost always satisfied in practice, for example, if each row of X is independently sampled from absolutely continuous distributions.

We are now ready to state the main lemma bounding the support size of π^*.

Lemma 3.4. *$\|\pi^*\|_0 \leq k + \frac{p(p+1)}{2}$ if Assumption 3.1 holds.*

Lemma 3.4 is established by an interesting observation into the properties of Karush-Kuhn-Tucker (KKT) conditions of the optimization problem Eq. (3), which involves a linear system with $\frac{p(p+1)}{2} + 1$ variables. We give the complete proof of Lemma 3.4 in Sec. 6.5.

Combining results from both Lemma 3.4 and Corollary 3.2 we arrive at the following theorem, which upper bounds the approximation ratio of the greedy removal procedure in Figure 1 initialized by the support of π^*.

Theorem 3.2. *Let π^* be the optimal solution of the without replacement version of Eq. (3) and $\hat{S}$ be the output of the greedy removal procedure in Figure 1 initialized with $S_0 = \{i \in [n] : \pi^*_i > 0\}$. If $k > p$ and Assumption 3.1 holds then*

$$F(\hat{S}; X) \leq \left(1 + \frac{p(p+1)}{2(k-p+1)}\right) F(S^*; X).$$

Under a slightly stronger condition that $k > 2p$, the approximation results in Theorem 3.2 can be simplified as $F(\hat{S}; X) \leq (1 + O(p^2/k))F(S^*; X)$. In addition, $F(\hat{S}; X) \leq (1 + o(1))F(S^*; X)$ if $p^2/k \to 0$, meaning that near-optimal experiment selection is achievable with computationally tractable methods if $O(k^2)$ design points are allowed in the selected subset.

3.4 Extensions

We discuss possible extension of our results beyond estimation of β_0 in the linear regression model.

3.4.1 Generalized linear models

In a generalized linear model $\mu(x) = \mathbb{E}[Y|x]$ satisfies $g(\mu(x)) = \eta = x^\top \beta_0$ for some known link function $g : \mathbb{R} \to \mathbb{R}$. Under regularity conditions (Van der Vaart, 2000), the maximum-likelihood estimator $\hat{\beta}_n \in \operatorname{argmax}_\beta \{\sum_{i=1}^n \log p(y_i|x_i;\beta)\}$ satisfies $\mathbb{E}\|\hat{\beta}_n - \beta_0\|_2^2 = (1+o(1))\mathrm{tr}(I(X,\beta_0)^{-1})$, where $I(X,\beta_0)$ is the Fisher's information matrix:

$$I(X,\beta_0) = -\sum_{i=1}^n \mathbb{E}\frac{\partial^2 \log p(y_i|x_i;\beta_0)}{\partial\beta\partial\beta^\top} = -\sum_{i=1}^n \left(\mathbb{E}\frac{\partial^2 \log p(y_i;\eta_i)}{\partial\eta_i^2}\right) x_i x_i^\top. \tag{5}$$

Here both expectations are taken over y conditioned on X and the last equality is due to the sufficiency of $\eta_i = x_i^\top \beta_0$. The experiment selection problem is then formulated to select a subset $S \subseteq [n]$ of size k, either with or without duplicates, that minimizes $\mathrm{tr}(I(X_S,\beta_0)^{-1})$.

It is clear from Eq. (5) that the optimal subset S^* depends on the unknown parameter β_0, which itself is to be estimated. This issue is known as the *design dependence* problem for generalized linear models (Khuri et al., 2006). On approach is to consider *locally optimal designs* (Khuri et al., 2006; Chernoff, 1953), where a consistent estimate $\check{\beta}$ of β_0 is first obtained on a initial design susbet and then $\check{\eta}_i = x_i^\top \check{\beta}$ is supplied to compute a more refined design subset to get the final estimate $\hat{\beta}$. With the initial estimate $\check{\beta}$ available, the optimal experiment selection rule becomes

$$S^* = \operatorname{argmin}_{|S|\le k} F_{\check{\beta}}(S;X) = \operatorname{argmin}_{|S|\le k} \mathrm{tr}\left[\left(\widetilde{X}_S^\top \widetilde{X}_S\right)^{-1}\right] = F(S;\widetilde{X}), \tag{6}$$

where $\widetilde{X} = (\widetilde{x}_1, \cdots, \widetilde{x}_n)^\top$ is defined as

$$\widetilde{x}_i = \sqrt{-\mathbb{E}\frac{\partial^2 \log p(y_i;\check{\eta}_i)}{\partial\eta^2}}\, x_i.$$

Note that under regularity conditions $-\mathbb{E}\frac{\partial^2 \log p(y_i;\check{\eta}_i)}{\partial\eta^2}$ is non-negative and hence the square-root is well-defined. Both results in Theorems 3.1 and 3.2 are valid with X replaced by $\widetilde{X}$ for generalized linear models. Below we consider two generalized linear model examples and derive explicit forms of $\widetilde{X}$.

Example 1: Logistic regression In a logistic regression model responses $y_i \in \{0,1\}$ are binary and the likelihood model is

$$p(y_i;\eta_i) = \psi(\eta_i)^{y_i}(1-\psi(\eta_i))^{1-y_i}, \qquad \text{where} \quad \psi(\eta_i) = \frac{e^{\eta_i}}{1+e^{\eta_i}}.$$

Simple algebra yields

$$\widetilde{x}_i = \sqrt{\frac{e^{\check{\eta}_i}}{(1+e^{\check{\eta}_i})^2}}\, x_i,$$

where $\check{\eta}_i = x_i^\top \check{\beta}$.

Example 2: Poisson count model In a Poisson count model the response variable y_i takes values of non-negative integers and follows a Poisson distribution with parameter $\lambda = e^{\eta_i} = e^{x_i^\top \beta_0}$. The likelihood model is formally defined as

$$p(y_i = r; \eta_i) = \frac{e^{\eta_i r} e^{-e^{\eta_i}}}{r!}, \qquad r = 0, 1, 2, \cdots .$$

Simple algebra yields

$$\widetilde{x}_i = \sqrt{e^{\check{\eta}_i}} x_i,$$

where $\check{\eta}_i = x_i^\top \check{\beta}$.

3.4.2 Delta's method

Suppose $g(\beta_0)$ is the quantity of interest, where $\beta_0 \in \mathbb{R}^p$ is the parameter in a linear regression model and $g : \mathbb{R}^p \to \mathbb{R}^m$ is some known function. Examples include the in-sample prediction $g(\beta) = X\beta$ and the out-sample prediction $g(\beta) = Z\beta$. Let $\hat{\beta}_n = (X^\top X)^{-1} X^\top y$ be the OLS estimate of β_0. If ∇g is continuously differentiable and $\hat{\beta}_n$ is consistent, then by the classical delta's method (Van der Vaart, 2000) $\mathbb{E}\|g(\hat{\beta}_n) - g(\beta_0)\|_2^2 = (1 + o(1))\sigma^2 \mathrm{tr}(\nabla g(\beta_0)(X^\top X)^{-1} \nabla g(\beta_0)^\top)$. The experiment selection problem of estimating $g(\beta_0)$ can then be formalized as

$$S^* = \mathrm{argmin}_{|S| \le k} F_{G_0}(S; X) = \mathrm{argmin}_{|S| \le k} \mathrm{tr}\left[G_0 \left(X_S^\top X_S \right)^{-1} \right],$$

where $G_0 = \nabla g(\beta_0)^\top \nabla g(\beta_0)$. If G_0 depends on the unknown parameter β_0 then the design dependence problem again exists, and a locally optimal solution can be obtained by replacing G_0 in the objective function with $\check{G} = \nabla g(\check{\beta})^\top \nabla g(\check{\beta})$ for some initial estimate $\check{\beta}$ of β_0.

If $\check{G}$ is invertible, then there exists invertible $p \times p$ matrix $\check{P}$ such that $\check{G} = \check{P}\check{P}^\top$ because $\check{G}$ is positive definite. The objective function can then be re-organized as

$$F_{\check{G}}(S; X) = \mathrm{tr}\left[\check{P}\check{P}^\top \left(X_S^\top X_S \right)^{-1} \right] = \mathrm{tr}\left[\left(\check{P}^{-1} X_S^\top X_S \check{P}^{-\top} \right)^{-1} \right] = F(S; X\check{P}^{-\top}),$$

where $F(S; X\check{P}^{-\top})$ is the ordinary A-optimality criterion define in Eq. (1). Our results in Theorems 3.1 and 3.2 remain valid by operating on the transformed matrix $X\check{P}^{-\top}$. In the case of $\check{G}$ being singular, it suffices to consider a regularized objective

$$F_{\check{G}_\lambda}(S; X) = \mathrm{tr}\left[\check{G}_\lambda \left(X_S^\top X_S \right)^{-1} \right], \qquad \text{where} \quad \check{G}_\lambda = \check{G} + \lambda I_{p \times p}.$$

Note that $\check{G}_\lambda$ is always invertible if $\lambda > 0$, because $\check{G} = \nabla g(\check{\beta})^\top \nabla g(\check{\beta})$ is positive semi-definite. Quality of such approximation can be bounded as

$$F_{\check{G}}(S; X) \le F_{\check{G}_\lambda}(S; X) \le F_{\check{G}}(S; X) + \lambda F(S; X).$$

Table 2: Isotropic design (i.i.d. t distribution with specificed degree of freedom (df)).

$F(S;X)$ without replacement $(\times 10^4)$; $\mathrm{df}=1$.							
k	L_1	L_2	L_3	L_4^*	L_5^*	$L_6^\dagger$	$f(\pi^*;X)$
$2p$	7.023(1.55)	.2099(.318)	1.726(.797)	1.215(.301)	.0388	.0388	.0377
$3p$	2.044(.417)	.0446(.007)	.3581(.150)	.3571(.131)	.0332	.0332	.0329
$4p$	.9636(.158)	.0364(.002)	.1385(.054)	.0595(.031)	.0309	.0309	.0308
$5p$	.5714(.129)	.0332(.001)	.0921(.029)	.0488(.015)	.0297	.0297	.0296
$6p$	.3963(.095)	.0313(.000)	.0645(.017)	.0437(.014)	.0289	.0289	.0289
$7p$	.2693(.049)	.0301(.000)	.0495(.011)	.0378(.001)	.0284	.0284	.0284
$8p$	.1944(.039)	.0293(.000)	.0399(.001)	.0382(.001)	.0281	.0281	.0281
$9p$	.1458(.026)	.0287(.000)	.0359(.005)	.0312(.004)	.0278	.0278	.0278
$10p$	.1199(.021)	.0283(.000)	.0337(.002)	.0309(.004)	.0277	.0277	.0277
$F(S;X)$ without replacement $(\times 10)$; $\mathrm{df}=3$.							
k	L_1	L_2	L_3	L_4^*	L_5^*	$L_6^\dagger$	$f(\pi^*;X)$
$2p$	6.048(.316)	4.568(.362)	5.150(.344)	2.203(.120)	1.634	1.632	1.531
$3p$	2.847(.124)	2.169(.116)	2.462(.132)	1.191(.040)	1.064	1.060	1.042
$4p$	1.843(.066)	1.424(.059)	1.630(.057)	.8767(.023)	.8184	.8174	.8108
$5p$	1.361(.043)	1.054(.031)	1.204(.037)	.7165(.012)	.6792	.6792	.6758
$6p$	1.088(.033)	.8464(.0235)	.9424(.0281)	.6126(.006)	.5877	.5878	.5861
$7p$	.8891(.022)	.7174(.019)	.7894(.026)	.5394(.006)	.5232	.5231	.5221
$8p$	.7574(.018)	.6158(.012)	.6718(.017)	.4872(.004)	.4745	.4745	.4740
$9p$	.6629(.014)	.5442(.009)	.5895(.011)	.4492(.003)	.4371	.4371	.4367
$10p$	.5866(.014)	.4903(.007)	.5291(.009)	.4165(.003)	.4068	.4068	.4065
$F(S;X)$ without replacement $(\times 10)$; $\mathrm{df}=5$.							
k	L_1	L_2	L_3	L_4^*	L_5^*	$L_6^\dagger$	$f(\pi^*;X)$
$2p$	5.248(.285)	4.753(.244)	5.030(.258)	3.136(.130)	2.303	2.285	1.982
$3p$	2.577(.080)	2.353(.063)	2.470(.082)	1.600(.036)	1.392	1.389	1.323
$4p$	1.698(.032)	1.557(.034)	1.634(.042)	1.112(.018)	1.023	1.022	1.000
$5p$	1.275(.024)	1.169(.025)	1.219(.022)	.8627(.009)	.8209	.8195	.8096
$6p$	1.019(.020)	.9311(.016)	.9739(.019)	.7184(.005)	.6903	.6897	.6851
$7p$	.8469(.012)	.7799(.012)	.8093(.014)	.6220(.004)	.5999	.5998	.5970
$8p$	.7217(.009)	.6698(.009)	.6958(.008)	.5486(.003)	.5332	.5331	.5315
$9p$	.6311(.007)	.5873(.007)	.6041(.008)	.4942(.002)	.4820	.4821	.4807
$10p$	.5604(.006)	.5254(.005)	.5395(.006)	.4498(.002)	.4410	.4410	.4402

*Methods proposed in this paper.

†The computationally inefficient Fedorov's exchange algorithm.

Table 3: Skewed design (transformed multivariate Gaussian with spectral decay $\lambda_j \propto j^{-\alpha}$.

$F(S;X)$ without replacement; $\alpha = 1$.							
k	L_1	L_2	L_3	L_4^*	L_5^*	$L_6^\dagger$	$f(\pi^*;X)$
$2p$	2.326(.095)	2.247(.109)	2.334(.106)	1.639(.059)	1.176	1.210	.9928
$3p$	1.162(.031)	1.117(.025)	1.150(.029)	.8273(.013)	.7103	.7078	.6620
$4p$	.7687(.016)	.7418(.016)	.7636(.014)	.5650(.006)	.5153	.5142	.4978
$5p$	.5782(.009)	.5585(.009)	.5737(.009)	.4357(.003)	.4084	.4078	.4005
$6p$	.4603(.008)	.4460(.006)	.4585(.008)	.3553(.002)	.3402	.3402	.3361
$7p$	.3842(.005)	.3708(.005)	.3816(.005)	.3029(.002)	.2927	.2926	.2905
$8p$	.3281(.004)	.3180(.004)	.3267(.003)	.2652(.001)	.2579	.2579	.2567
$9p$	.2882(.003)	.2797(.003)	.2864(.002)	.2372(.001)	.2315	.2314	.2307
$10p$	.2560(.002)	.2478(.002)	.2544(.002)	.2151(.001)	.2105	.2105	.2099
$F(S;X)$ without replacement $(\times 10^3)$; $\alpha = 3$.							
k	L_1	L_2	L_3	L_4^*	L_5^*	$L_6^\dagger$	$f(\pi^*;X)$
$2p$	.7951(.040)	.7527(.038)	.8013(.038)	.4837(.016)	.3624	.3554	.3087
$3p$	.3943(.012)	.3785(.014)	.3935(.013)	.2504(.004)	.2181	.2178	.2063
$4p$	.2629(.007)	.2518(.007)	.2643(.006)	.1736(.002)	.1600	.1597	.1560
$5p$	.1954(.004)	.1900(.004)	.1968(.003)	.1359(.001)	.1284	.1283	.1263
$6p$	.1568(.003)	.1506(.003)	.1571(.003)	.1129(.001)	.1080	.1078	.1068
$7p$	.1308(.002)	.1264(.002)	.1311(.002)	.0969(.000)	.0936	.0935	.0930
$8p$	.1120(.002)	.1084(.002)	.1127(.001)	.0855(.000)	.0831	.0830	.0827
$9p$	.0979(.001)	.0947(.001)	.0983(.001)	.0767(.000)	.0749	.0749	.0749
$10p$	.0873(.001)	.0845(.001)	.0872(.001)	.0700(.000)	.0685	.0685	.0683
$F(S;X)$ without replacement $(\times 10^6)$; $\alpha = 5$.							
k	L_1	L_2	L_3	L_4^*	L_5^*	$L_6^\dagger$	$f(\pi^*;X)$
$2p$	1.165(.074)	1.105(.072)	1.152(.081)	.6878(.027)	.5020	.5093	.4337
$3p$	.5812(.024)	.5496(.022)	.5713(.024)	.3496(.008)	.3054	.3077	.2896
$4p$	.3831(.013)	.3694(.012)	.3792(.011)	.2443(.003)	.2251	.2243	.2188
$5p$	.2877(.008)	.2762(.006)	.2843(.006)	.1911(.002)	.1801	.1800	.1774
$6p$	.2282(.005)	.2206(.004)	.2274(.004)	.1588(.001)	.1518	.1517	.1503
$7p$	.1907(.004)	.1842(.003)	.1884(.004)	.1368(.001)	.1319	.1319	.1311
$8p$	.1637(.003)	.1587(.002)	.1617(.003)	.1208(.001)	.1173	.1173	.1168
$9p$	.1436(.002)	.1387(.002)	.1415(.002)	.1088(.001)	.1060	.1060	.1057
$10p$	.1272(.002)	.1236(.002)	.1258(.001)	.0993(.000)	.0972	.0972	.0970

*Methods proposed in this paper.

†The computationally inefficient Fedorov's exchange algorithm.

Table 4: Running time (seconds) / no. of iterations for L_5^* and $L_6^\dagger$.

Isotropic design (t distribution), $\mathrm{df} = 1$.					
$k =$	$2p$	$4p$	$6p$	$8p$	$10p$
L_5^*	1.8/31	0.8/19	1.7/26	0.9/14	0.3/9
$L_6^\dagger$	140.1/90	425.9/156	767.5/216	993.7/245	1077/253
Isotropic design (t distribution), $\mathrm{df} = 3$.					
$k =$	$2p$	$4p$	$6p$	$8p$	$10p$
L_5^*	0.6/14	0.4/8	0.3/7	0.2/5	0.2/5
$L_6^\dagger$	158.3/104	438.0/161	802.5/223	985.2/242	1105/252
Skewed design (transformed Gaussian), $\alpha = 5$.					
$k =$	$2p$	$4p$	$6p$	$8p$	$10p$
L_5^*	.8/16	.7/12	.5/9	.4/8	.4/8
$L_6^\dagger$	182.9/120	487.1/180	753.4/212	935.8/230	1057/250

4 Simulations

We report selection performance (measured in terms of $F(S; X)$) on synthetic data. Only the without replacement model is considered. In all simulations, number of experimental pool n is set to 1000, number of variables p is set to 50, number of selected design points k ranges from $2p$ to $10p$. For randomized methods, we run them for 20 independent trials under each setting and report the median as well as standard deviation (in brackets). Though Fedorov's exchange algorithm is randomized in nature (outcome depending on the initialization), we only run it once for each setting because of its expensive computational requirement. It is observed that in practice, the algorithm's outcome is stable and not sensitive to initializations.

4.1 Data generation

The design pool X is generated so that each row of X is sampled from some underlying distribution. We consider two types of distributions for generating X. The first type of distribution is roughly isotropic, in the sense that each variable of X is independent and of the same magnitude. For this setting, we use t distribution as the underlying distribution for generating each entry in X, with degrees of freedom ranging from 1 to 10. Note that for a t distribution of df 1, the mean of each entry does not exist, which tests the robustness of experiment selection methods. Similar settings were also considered in previous work (Ma et al., 2015) for subsampling purposes.

The second setting is skewed distributions, where correlation between variables are introduced and some variables (or linear combination of variables) have significantly higher variance on the others. To generate such design pool X, each row of X is sampled i.i.d. from a multivariate Gaussian distribution $\mathcal{N}_p(0, \Sigma_0)$ with $\Sigma_0 = U\Lambda U^\top$, where U is a random orthogonal matrix and $\Lambda = \mathrm{diag}(\lambda_1, \cdots, \lambda_p)$ controls the skewedness or condition of X. A power-law decay of λ, $\lambda_j = j^{-\alpha}$, is imposed, with small α corresponding to "flat" distribution and large α corresponding to "skewed" distribution.

4.2 Methods

The methods that we compare are listed below:

- L_1 (uniform sampling): each row of X is sampled uniformly at random, without replacement.
- L_2 (leverage score sampling): each row of X, x_i, is sampled without replacement, with probability proportional to its leverage score $x_i^\top (X^\top X)^{-1} x_i$. This strategy is considered in Ma et al. (2015) for subsampling in linear regression models.
- L_3 (predictive length sampling): each row of X, x_i, is sampled without replacement, with probability proportional to its ℓ_2 norm $\|x_i\|_2$. This strategy is derived in Zhu et al. (2015).
- L_4^* (sampling based selection with π^*): the sampling based method that is described in Sec. 3.2, based on π^*, the optimal solution of Eq. (3).
- L_5^* (greedy based selection with π^*): the greedy based method that is described in Sec. 3.3.
- $L_6^\dagger$ (Fedorov's exchange algorithm): the Fedorov's exchange algorithm (Miller & Nguyen, 1994) with the A-optimality objective. Note that this method does not have guarantees, nor does it provably terminate in polynomial number of exchanges. Description of the algorithm is placed in Appendix C.

4.3 Performance

Objective values $F(\hat{S}; X) = \mathrm{tr}((X_{\hat{S}}^\top X_{\hat{S}})^{-1})$ are reported in Tables 2 and 3. We also report $f(\pi^*; X)$, the optimal value of the continuously relaxed objective in Eq. (3), which is technically *not* a feasible experiment selection solution but serves as a lower bound, thanks to Fact 3.1. We observe that both L_4^* and L_5^* outperform existing sampling based approaches in the literature (L_1 to L_4) by a large margin, and their performance is comparable with the performance of Fedorov's exchange algorithm, which is computationally much slower (cf. Table 4). Another interesting fact is the lower bound given by $f(\pi^*; X)$: under most settings the lower bound matches quite well with the experiment subset selected by L_5^* and $L_6^\dagger$, which suggests that the continuous relaxation Eq. (3) is appropriate for practical applications.

Table 4 reports the running time and number of iterations of L_5^* (greedy based selection) and $L_6^\dagger$ (Fedorov's exchange algorithm). In general, the greedy method is 100 to 1000 times faster than the exchange algorithm, and also converges in much fewer iterations. In addition, the continuous optimization problem Eq. (3) seems easier to solve (i.e., converges faster) when the number of desired experiments k is large.

5 Discussion

We discuss on potential improvement of analysis presented in this paper.

5.1 Sampling based method

To fully understand the finite-sample behavior of the sampling method introduced in Sec. 3.2, it is instructive to relate it to the *graph spectral sparsification* problem (Spielman & Srivastava, 2011) in theoretical computer science: Given a directed weighted graph $G = (V, E, W)$, find a subset $\tilde{E} \subseteq E$ and new weights $\tilde{W}$ such that $\tilde{G} = (V, \tilde{E}, \tilde{W})$ is a *spectral sparsification* of G, which means there exists $\epsilon \in (0, 1)$ such that for any vector $z \in \mathbb{R}^{|V|}$:

$$(1-\epsilon)\sum_{(u,v)\in E} w_{uv}(z_u - z_v)^2 \leq \sum_{(u,v)\in \tilde{E}} \tilde{w}_{uv}(z_u - z_v)^2 \leq (1+\epsilon)\sum_{u,v\in E} w_{uv}(z_u - z_v)^2.$$

Define $B_G \in \mathbb{R}^{|E|\times|V|}$ to be the *signed edge-vertex incidence matrix*, where each row of B_G corresponds to an edge in E, each column of B_G corresponds to a vertex in G, and $[B_G]_{ij} = 1$ if vertex j is the head of edge i, $[B_G]_{ij} = -1$ if vertex j is the tail of edge i, and $[B_G]_{ij} = 0$ otherwise. The spectral sparsification requirement can then be equivalently written as

$$(1-\epsilon)z^\top (B_G^\top W B_G) z \leq z^\top (B_{\tilde{G}}^\top \tilde{W} B_{\tilde{G}}) z \leq (1+\epsilon) z^\top (B_G^\top W B_G) z.$$

The similarity of graph sparsification and the experimental selection problem is clear: $B_G \in \mathbb{R}^{n\times p}$ would be the known pool of design points and $W = \mathrm{diag}(\pi^*)$ is a diagonal matrix with the optimal weights π^* obtained by solving Eq. (3). The objective is to seek a small subset of rows in B_G (i.e., the sparsified edge set $\tilde{E}$) which is a spectral approximation of the original $B_G^\top \mathrm{diag}(\pi^*) B_G$. Approximation of the A-optimality criterion or any other eigen-related quantity immediately follows. One difference is that in linear regression each row of B_G is no longer a $\{\pm 1, 0\}$ vector. However, we consider this to be a minor difference as it does not intefere with the spectral properties of B_G.

The spectral sparsification problem where $\tilde{W}$ can be arbitrarily designed is completely solved (Spielman & Srivastava, 2011; Batson et al., 2012), where the size of the selected edge subset is allowed to be *linear* to the number of vertices, or in terminology of our problem, $k \asymp p$. Unfortunately, both methods require the power of arbitrary designing the weights in $\tilde{W}$, which is generally not available in experiment selection problems (i.e., cannot set noise variance or signal strength arbitrarily for individual design points). Recently, it was proved that when the original graph is unweighted ($W = I$), it is also possible to find unweighted linear-sized edge sparsifiers ($\tilde{W} \propto I$) (Marcus et al., 2015a,b; Anderson et al., 2014). This remarkable result leads to the solution of the long-standing Kardison-Singer problem. However, the condition that the *original* weights W are uniform is not satisfied in the linear regression problem, where the optimal solution π^* may be far from uniform. The experiment selection problem somehow falls in between, where an *unweighted* sparsifier is desired for a *weighted* graph. This leads us to the following question:

Question 5.1. *Given a* weighted *graph $G = (V, E, W)$, under what conditions are there small edge subset $\tilde{E} \subseteq E$ with* uniform *weights $\tilde{W} \propto I$ such that $\tilde{G} = (V, \tilde{E}, \tilde{W})$ is spectral approximation of G?*

The answer to the above question, especially the smallest possible edge size $|\tilde{E}|$, would have immediate consequences on finite-sample conditions of k and p in the experiment selection problem.

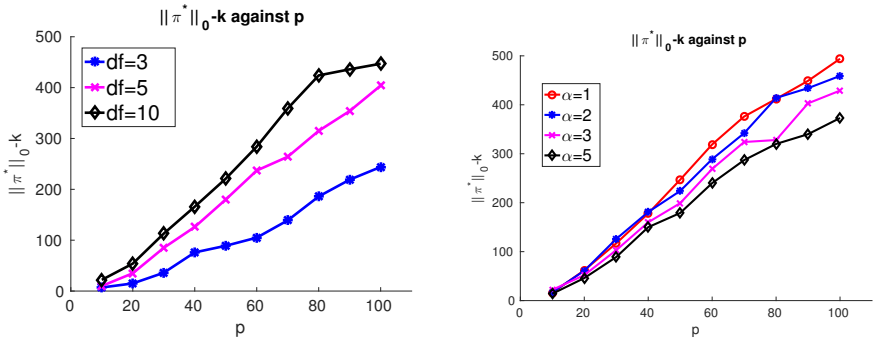

Figure 1: Empirical verification on the rate of $\|\pi^*\|_0$. Left: isotropic design (t distribution); Right: skewed design (transformed multivariate Gaussian) with spectral decay $\lambda_j \propto j^{-\alpha}$.

5.2 Greedy based method

Corollary 3.2 shows that the approximation quality of the greedy based method depends crucially upon $\|\pi^*\|_0$, the support size of the optimal solution π^*. In Lemma 3.4 we formally established that $\|\pi^*\|_0 \le k + p(p+1)/2$ under mild conditions; however, we conjecture the $p(p+1)/2$ term is loose and could be improved to $O(p)$ in general cases.

In Fig 1 we plot $\|\pi^*\|_0 - k$ against number of variables p, where p ranges from 10 to 100 and k is set to $3p$. The other simulation settings are kept unchanged. We observe that in all settings $\|\pi^*\|_0 - k$ scales linearly with p, suggesting that $\|\pi^*\|_0 \le k + O(p)$. Furthermore, the slope of the scalings does not seem to depend on the conditioning of X, as shown in the right panel of Fig 1 where the conditioning of X is controlled by the spectral decay rate α. This is in constrast to the analysis of the sampling based method (Theorem 3.1), in which the finite sample bound depends crucially upon the conditioning of X.

6 Proofs

6.1 Proof of facts in Sec. 3.1

Proof of Fact 3.1. Define $\pi = (\pi_1, \cdots, \pi_n)$ as

$$\pi_i = \#\{i \in S^*\}, \qquad i = 1, \cdots, n.$$

It is easy to verify that $\pi \ge 0$, $\|\pi\|_1 = k$ and in addition $\|\pi\|_\infty \le 1$ for without replacement models, because duplicates of rows are disallowed. In addition,

$$f(\pi; X) = \mathrm{tr}\left[\left(\sum_{i \in S^*} \pi_i x_i x_i^\top\right)^{-1}\right] = \mathrm{tr}\left[\left(X_{S^*}^\top X_{S^*}\right)^{-1}\right] = F(S^*; X).$$

Subsequently, $f(\pi^*; X) \le f(\pi; X) = F(S^*; X)$. □

Proof of Fact 3.2. Let $A = X^\top \mathrm{diag}(\pi) X$ and $B = X^\top \mathrm{diag}(\pi') X$. By definition, $B = A + \Delta$ where $\Delta = X^\top \mathrm{diag}(\pi' - \pi) X$ is positive semi-definite. Subsequently, $\sigma_\ell(B) \ge \sigma_\ell(A)$ for all

$\ell = 1, \cdots, p$. We then have that

$$f(\pi'; X) = \mathrm{tr}(B^{-1}) = \sum_{\ell=1}^{p} \sigma_\ell(B)^{-1} \leq \sum_{\ell=1}^{p} \sigma_\ell(A)^{-1} = \mathrm{tr}(A^{-1}) = f(\pi; X).$$

Note also that if $\pi' \neq \pi$ then $\Delta \neq 0$ and hence there exists at least one ℓ with $\sigma_\ell(B) < \sigma_\ell(A)$. Therefore, the equality holds if and only if $\pi = \pi'$. □

Proof of Fact 3.3. Suppose $\|\pi^*\|_1 < k$. Then there exists some coordinate $j \in \{1, \cdots, n\}$ such that $\pi^*_j < 1$. Define π' as $\pi'_i = \pi^*_i$ for $i \neq j$ and $\pi'_j = \min\{1, \pi^*_j + k - \|\pi^*\|_1\}$. Then π' is also a feasible solution. On the other hand, By Fact 3.2 we have that $f(\pi'; X) < f(\pi^*; X)$, contradicting the optimality of π^*. Therefore, $\|\pi^*\|_1 = k$. □

6.2 Proof of Lemma 3.1

Define $\Phi = \mathrm{diag}(\pi^*)$ and $\Pi = \Phi^{1/2} X \Sigma_*^{-1} X^\top \Phi^{1/2} \in \mathbb{R}^{n \times n}$. The following proposition lists properties of Π:

Proposition 6.1 (Properties of projection matrix)**.** *The following properties for* Π *hold:*

1. Π *is a projection matrix. That is,* $\Pi^2 = \Pi$.
2. $\mathrm{Range}(\Pi) = \mathrm{Range}(\Phi^{1/2} X)$.
3. *The eigenvalues of* Π *are 1 with multiplicity* p *and 0 with multiplicity* $n - p$.
4. $\Pi_{ii} = \|\Pi_{i,\cdot}\|_2^2 = \pi^*_i x_i^\top \Sigma_*^{-1} x_i$.

Proof. Proof of 1: By definition, $\Sigma_* = X^\top \Phi X$ and subsequently

$$\begin{aligned} \Pi^2 &= \Phi^{1/2} X \Sigma_*^{-1} X^\top \Phi^{1/2} \Phi^{1/2} X \Sigma_*^{-1} X^\top \Phi^{1/2} \\ &= \Phi^{1/2} X (X^\top \Phi X)^{-1} X^\top \Phi X (X^\top \Phi X)^{-1} \Phi^{1/2} \\ &= \Phi^{1/2} X (X^\top \Phi X)^{-1} X^\top \Phi^{1/2} = \Pi. \end{aligned}$$

Proof of 2: First note that $\mathrm{Range}(\Pi) = \mathrm{Range}(\Phi^{1/2} X \Sigma_*^{-1} X^\top \Phi^{1/2}) \subseteq \mathrm{Range}(\Phi^{1/2} X)$. For the other direction, take arbirary $u \in \mathrm{Range}(\Phi^{1/2} X)$ and express u as $u = \Phi^{1/2} X v$ for some $v \in \mathbb{R}^p$. We then have

$$\begin{aligned} \Pi u &= \Phi^{1/2} X \Sigma_*^{-1} X^\top \Phi^{1/2} u \\ &= \Phi^{1/2} X (X^\top \Phi X)^{-1} X^\top \Phi^{1/2} \Phi^{1/2} X v \\ &= \Phi^{1/2} X v = u \end{aligned}$$

and hence $u \in \mathrm{Range}(\Pi)$.

Proof of 3: Because $\Sigma_* = X^\top \Phi X$ is invertible, the $n \times p$ matrix $\Phi^{1/2} X$ must have full column rank and hence $\ker(\Phi^{1/2} X) = \{0\}$. Consequently, $\dim(\mathrm{Range}(\Pi)) = \dim(\mathrm{Range}(\Phi^{1/2} X)) = p - \dim(\ker(\Phi^{1/2} X)) = p$. On the other hand, the eigenvalues of Π must be either 0 or 1 because Π is a projection matrix. So the eigenvalues of Π are 1 with multiplicity p and 0 with multiplicity $n - p$.

Proof of 4: By definition,

$$\Pi_{ii} = \sqrt{\pi_i^*} x_i^\top \Sigma_*^{-1} x_i \sqrt{\pi_i^*} = \pi_i^* x_i^\top \Sigma_*^{-1} x_i.$$

In addition, Π is a symmetric projection matrix. Therefore,

$$\Pi_{ii} = [\Pi^2]_{ii} = \Pi_{i,\cdot}^\top \Pi_{i,\cdot} = \|\Pi_{i,\cdot}\|_2^2.$$

□

The following lemma shows that a spectral norm bound over deviation of the projection matrix implies spectral approximation of the underlying (weighted) covariance matrix.

Lemma 6.1 (Spielman & Srivastava (2011), Lemma 4)**.** *Let $\Pi = \Phi^{1/2} X \Sigma_*^{-1} X^\top \Phi^{1/2}$ and W be an $n \times n$ non-negative diagonal matrix. If $\|\Pi W \Pi - \Pi\|_2 \le \epsilon$ for some $\epsilon \in (0, 1/2)$ then*

$$(1-\epsilon) u^\top \Sigma_* u \le u^\top \widetilde{\Sigma}_* u \le (1+\epsilon) u^\top \Sigma_* u, \quad \forall u \in \mathbb{R}^p,$$

where $\Sigma_* = X^\top \Phi X$ and $\widetilde{\Sigma}_* = X^\top W^{1/2} \Phi W^{1/2} X$.

Define $\Sigma_* = X^\top \Phi X$ and $\widetilde{\Sigma}_* = X_{\hat{S}}^\top X_{\hat{S}}$, where $\hat{S}$ is a size-k subset of $[n]$ with elements sampled i.i.d. with replacement from the categorical distribution specified by π^*/k. $\widetilde{\Sigma}_*$ can be written as $\widetilde{\Sigma}_* = X^\top W^{1/2} \Phi W^{1/2} X$ where W is a diagonal matrix with $W_{ii} = \frac{\#\{i \in \hat{S}\}}{\sqrt{\pi_i^*}}$. We then have the following lemma bounding the spectral norm perturbation between $\Pi W \Pi$ and Π:

Lemma 6.2. *Let $W_{ii} = \frac{\#\{i \in \hat{S}\}}{\sqrt{\pi_i^*}}$. For any $\delta \in (0, 1)$*

$$\Pr\left[\|\Pi W \Pi - \Pi\|_2 > \delta\right] \le 2 \exp\left\{ - \frac{\delta^2}{\|X\|_\infty^2 \|\Sigma_*^{-1}\|_2 \log k} \right\}.$$

Proof. Define n-dimensional random vector v as

$$\Pr\left[v = \sqrt{\frac{k}{\pi_i^*}} \Pi_{i,\cdot} \right] = \frac{\pi_i^*}{k}, \qquad i = 1, \cdots, n.$$

Let $v_1, \cdots, v_k$ be i.i.d. copies of v. It is easy to verify that $\Pi W \Pi$ is equally distributed with $\bar{V} = \frac{1}{k} \sum_{\ell=1}^k v_\ell v_\ell^\top$. In addition,

$$\mathbb{E} \bar{V} = \sum_{i=1}^k \Pi_{i,\cdot} \Pi_{i,\cdot}^\top = \Pi$$

and

$$\|v\|_2^2 = \frac{k}{\pi_i^*} \|\Pi_{i,\cdot}\|_2^2 = \frac{k}{\pi_i^*} \pi_i^* x_i^\top \Sigma_*^{-1} x_i \le k \|X\|_\infty^2 \|\Sigma_*^{-1}\|_2$$

almost surely. Also note that $\|\Pi\|_2 = 1$ because Π is a projection matrix. Apply Lemma A.2 with $M^2 = k \|X\|_\infty^2 \|\Sigma_*^{-1}\|_2$:

$$\Pr\left[\|\bar{V} - \Pi\|_2 > \delta\right] \le 2 \exp\left\{ - \frac{\delta^2}{\|X\|_\infty^2 \|\Sigma_*^{-1}\|_2 \log k} \right\}.$$

□

Under the condition that $\|X\|_\infty^2 \|\Sigma_*^{-1}\|_2 \log k \to 0$, with probability $1 - o(1)$ the inequality $\|\Pi W \Pi - \Pi\|_2 = o(1)$ holds and hence $u^\top \tilde{\Sigma}_* u = (1 + o(1)) u^\top \Sigma_* u$ holds uniformly for all $u \in \mathbb{R}^p$ with probability $1 - o(1)$. The proof of Lemma 3.1 is then completed.

6.3 Proof of Lemma 3.2

Define $M = \max_{1 \le i \le n} \#\{i \in \hat{S}\}$ and let $M_0 \in \{0, 1, \cdots, k-1\}$ be an arbitrary threshold. Suppose $\hat{S} = \{I_1, \cdots, I_k\}$, where $I_1, \cdots, I_k$ are i.i.d. sampled with replacement from the categorical distribution specified by π^*/k. Because $\|\pi^*\|_\infty \le 1$, for any distinct $j_1, j_2, \cdots, j_{M_0+1} \in [k]$ it holds that $\Pr[I_{j_1} = \cdots = I_{j_{M_0+1}}] \le k^{-M_0}$. By union bound we have

$$\begin{aligned}\Pr[M > M_0] &\le \sum_{1 \le j_1 < \cdots < j_{M_0+1} \le k} \Pr[I_{j_1} = I_{j_2} = \cdots = I_{j_{M_0+1}}] \\ &\le \binom{k}{M_0+1} \left(\frac{1}{k}\right)^{M_0} = k \binom{k}{M_0+1} \left(\frac{1}{k}\right)^{M_0+1}.\end{aligned}$$

Let $Z_1, \cdots, Z_k$ be i.i.d. distributed binary random variables with $\Pr[Z_j = 1] = 1/k$ and $\Pr[Z_j = 0] = 1 - 1/k$. By binary theorem,

$$\begin{aligned}\Pr\left[Z_1 + \cdots + Z_k > M_0\right] &= \sum_{m=M_0+1} k \binom{k}{m} \left(\frac{1}{k}\right)^m \left(1 - \frac{1}{k}\right)^{k-m} \\ &\ge \binom{k}{M_0+1} \left(\frac{1}{k}\right)^{M_0+1} \left(1 - \frac{1}{k}\right)^{k-M_0-1} \\ &\ge \left(1 - \frac{1}{k}\right)^k \cdot \binom{k}{M_0+1} \left(\frac{1}{k}\right)^{M_0+1} \\ &\ge \frac{1}{e} \cdot \binom{k}{M_0+1} \left(\frac{1}{k}\right)^{M_0+1}. \end{aligned} \tag{7}$$

On the other hand, applying Chernoff bound (Lemma A.3) with $\mu = \mathbb{E}[Z_1 + \cdots + Z_k] = 1$ we have

$$\Pr\left[Z_1 + \cdots + Z_k > M_0\right] = \Pr\left[Z_1 + \cdots + Z_k \ge M_0 + 1\right] \le e^{-M_0/3}. \tag{8}$$

Combining Eq. (7,8) with we arrive at

$$\Pr[M > M_0] \le k \exp\left(-\frac{M_0}{3} + 1\right).$$

Therefore, $M = O_P(\log k)$, which is to be demonstrated.

6.4 Proof of Theorem 3.1

The with replacement part of Theorem 3.1 is immediate from Lemma 3.1. We focus on the without replacement part in the rest of the proof. Let $\hat{S} = \{I_1, I_2, \cdots, I_k\}$, where $I_1, \cdots, I_k$ are sampled *without* replacement. An equivalent construction of $\hat{S}$ is follows: for $\ell = 1, 2, \cdots$, sample I'_ℓ *with* replacement; if I'_ℓ is already in $\hat{S}$ then discard this sample; otherwise add I'_ℓ to $\hat{S}$. Repeat this

procedure until $|\hat{S}| = k$. Suppose $I'_1, \cdots, I'_k$ are the first k samples obtained by sampling with replacement and let $\hat{S}' = \{I'_1, \cdots, I'_k\}$ be the multi-set and $\check{S}'$ be the set of distinct elements in $\hat{S}'$. Let M denote the maximum number of duplicates in $\hat{S}'$. We then have that

$$F(\hat{S}'; X) = \mathrm{tr}\left[\left(\sum_{i=1}^{n} \#\{i \in \hat{S}'\} x_i x_i^\top\right)^{-1}\right] = \frac{1}{M}\mathrm{tr}\left[\left(\sum_{i=1}^{n} \frac{\#\{i \in \hat{S}'\}}{M} x_i x_i^\top\right)^{-1}\right]$$
$$\geq \frac{1}{M}\mathrm{tr}\left[\left(\sum_{i \in \check{S}'} x_i x_i^\top\right)^{-1}\right] = \frac{1}{M} \cdot F(\check{S}'; X).$$

On the other hand, because $\check{S}' \subseteq \hat{S}$ we have that $F(\check{S}'; X) \geq F(\hat{S}; X)$. Subsequently,

$$F(\hat{S}; X) \leq M \cdot F(\hat{S}'; X).$$

By Lemma 3.1 and 3.2 we know that $M = O_P(\log k)$ and $F(\hat{S}'; X) = (1 + o(1))F(S^*; X)$. Consequently, we have that $F(\hat{S}; X) \leq O(\log k) \cdot F(S^*; X)$, which is to be demonstrated.

6.5 Proof of Lemma 3.4

Let $f(\pi; \lambda, \tilde{\lambda}, \mu)$ be the Lagrangian muliplier function of the without replacement formulation of Eq. (3):

$$f(\pi; \lambda, \tilde{\lambda}, \mu) = f(\pi; X) - \sum_{i=1}^{n} \lambda_i \pi_i + \sum_{i=1}^{n} \tilde{\lambda}_i \left(\pi_i - \frac{1}{k}\right) + \mu\left(\sum_{i=1}^{n} \pi_i - 1\right).$$

Here $\{\lambda_i\}_{i=1}^n \geq 0$, $\{\tilde{\lambda}_i\}_{i=1}^n \geq 0$ and $\mu \geq 0$ are Lagrangian multipliers for constraints $\pi_i \geq 0$, $\pi_i \leq 1$ and $\sum_i \pi_i \leq k$, respectively. By KKT condition, $\frac{\partial f}{\partial \pi_i}\big|_{\pi^*} = 0$ and hence

$$-\frac{\partial f}{\partial \pi_i}\bigg|_{\pi^*} = x_i^\top \Sigma_*^{-2} x_i = \tilde{\lambda}_i - \lambda_i + \mu, \qquad i = 1, \cdots, n,$$

where $\Sigma_* = X^\top \mathrm{diag}(\pi^*) X$ is a $p \times p$ positive definite matrix.

Split the index set $[n]$ into three disjoint sets defined as $A = \{i \in [n] : \pi_i^* = 1\}$, $B = \{i \in [n] : 0 < \pi_i^* < 1\}$ and $C = \{i \in [n] : \pi_i^* = 0\}$. Note that $\|\pi^*\|_0 = |A| + |B|$ and $|A| \leq k$. Therefore, to upper bound $\|\pi^*\|_0$ it suffices to upper bound $|B|$. By complementary slackness, for all $i \in B$ we have that $\tilde{\lambda}_i = \lambda_i = 0$; that is,

$$x_i^\top \Sigma_*^{-2} x_i = \langle \phi(x_i), \psi(\Sigma_*^{-2}) \rangle = \mu, \qquad \forall i \in B, \tag{9}$$

where $\phi : \mathbb{R}^p \to \mathbb{R}^{p(p+1)/2}$ is the mapping defined in Assumption 3.1 and $\psi(\cdot)$ takes the upper triangle of a symmetric matrix and vectorizes it into a $\frac{p(p-1)}{2}$-dimensional vector. Assume by way of contradiction that $|B| > p(p+1)/2$ and let $x_1, \cdots, x_{p(p+1)/2+1}$ be arbitrary distinct $\frac{p(p+1)}{2} + 1$ rows whose indices belong to B. Eq. (9) can then be cast as a homogenous linear system with

$\frac{p(p+1)}{2}+1$ variables and equations as follows:

$$\begin{bmatrix} \tilde{\phi}(x_1) \\ \tilde{\phi}(x_2) \\ \vdots \\ \tilde{\phi}(x_{p(p+1)/2+1}) \end{bmatrix} \begin{bmatrix} \psi(\Sigma_*^{-2}) \\ -\mu \end{bmatrix} = 0.$$

Under Assumption 3.1, $\tilde{\Phi} = [\tilde{\phi}(x_1); \cdots; \tilde{\phi}(x_{p(p+1)/2+1})]^\top$ is invertible and hence both $\psi(\Sigma_*^{-2})$ and μ must be zero. This contradicts the fact that Σ_*^{-2} is positive definite.

A Technical lemmas

Lemma A.1. *Let $\mathbb{B} = \{x \in \mathbb{R}^p : \|x\|_2^2 \le B^2\}$ and $\mathrm{vol}(\mathbb{B}) = \int_{\mathbb{B}} 1 \mathrm{d}x$ be the volume of $\mathbb{B}$. Then $\int_{\mathbb{B}} xx^\top \mathrm{d}x = \frac{B^2}{p+2}\mathrm{vol}(\mathbb{B}) I_{p\times p}$.*

Proof. Let U be the uniform distribution in the p-dimensional ball of radius B. By definition, $\int_{\mathbb{B}} xx^\top \mathrm{d}x = \mathrm{vol}(\mathbb{B})\mathbb{E}_{x\sim U}[xx^\top]$. By symmetry, $\mathbb{E}_{x\sim U}[xx^\top] = c \cdot I_{p\times p}$ for some constant c that depends on p and B. To determine the constant c, note that

$$c = \mathbb{E}[x_1^2] = \frac{1}{p}\mathbb{E}\|x\|_2^2 = \frac{1}{p}\frac{\int_0^B r^{p-1}\cdot r^2 \mathrm{d}r}{\int_0^B r^{p-1}\mathrm{d}r} = \frac{B^2}{p+2}.$$

□

Lemma A.2 (Rudelson & Vershynin (2007))**.** *Let x be a p-dimensional random vector such that $\|x\|_2 \le M$ almost surely and $\|\mathbb{E}xx^\top\|_2 \le 1$. Let $x_1, \cdots, x_n$ be i.i.d. copies of x. Then for every $t \in (0,1)$*

$$\Pr\left[\left\|\frac{1}{n}\sum_{i=1}^n x_i x_i^\top - \mathbb{E}xx^\top\right\|_2 > t\right] \le 2\exp\left\{-C\cdot\frac{nt^2}{M^2\log n}\right\},$$

where $C > 0$ is some universal constant.

Lemma A.3 (Multiplicative Chernoff bound)**.** *Suppose $X_1, \cdots, X_n$ are independent random variables taking values in $\{0,1\}$. Let $X = X_1 + \cdots + X_n$ and $\mu = \mathbb{E}X$. Then for any $\delta > 0$,*

$$\Pr[X \ge (1+\delta)\mu] \le \begin{cases} \exp\left\{-\frac{\delta^2\mu}{3}\right\}, & 0 < \delta < 1; \\ \exp\left\{-\frac{\delta\mu}{3}\right\}, & \delta \ge 1. \end{cases}$$

B Optimization methods

Two algorithms for optimizing Eq. (3) are described. The SDP formulation is of theoretical interest only and the projected gradient descent algorithm is practical, which also enjoys theoretical convergence guarantees.

SDP formulation For $\pi \in \mathbb{R}^n$ define $A(\pi) = \sum_{i=1}^n \pi_i x_i x_i^\top$, which is a $p \times p$ positive semidefinite matrix. By definition, $f(\pi; X) = \sum_{j=1}^p e_j^\top A(\pi)^{-1} e_j$, where e_j is the p-dimensional vector with only pth coordinate being 1. Subsequently, Eq. (3) is equivalent to the following SDP problem:

$$\min_{\pi,t} \sum_{j=1}^p t_j \quad \text{subject to} \quad 0 \le \pi_i \le 1, \ \sum_{i=1}^n \pi_i \le k, \ \ \text{diag}(B_1, \cdots, B_p) \succeq 0,$$

where

$$B_j = \begin{bmatrix} A(\pi) & e_j \\ e_j^\top & t_j \end{bmatrix}, \qquad j = 1, \cdots, p.$$

Global optimal solution of an SDP can be computed in polynomial time (Vandenberghe & Boyd, 1996). However, this formulation is not intended for practical computation because of the large number of variables in the SDP system. First-order methods such as projected gradient descent is a more appropriate choice for practical computation.

Projected gradient descent For any convex set S and point x let $\mathcal{P}_S(x) = \text{argmin}_{y \in S} \|x - y\|_2$ denote the ℓ_2 projection of x onto S. The projected gradient descent algorithm is a general purpose method to solve convex constrained smooth convex optimization problems of the form

$$\min_x f(x) \quad \text{subject to} \quad x \in S.$$

The algorithm (with step size selected via backtracking line search) iterates until desired optimization accuracy is reached:

input : backtracking parameters $\alpha \in (0, 1/2]$, $\beta \in (0, 1)$.
output: $\hat{x}$, approximate solution of the optimization problem.
Initialization: x_0, $t = 0$.
1. Compute gradient $g_t = \nabla f(x_t)$.
2. Find the smallest integer $s \ge 0$ such that $f(x') - f(x_t) \le \alpha g_t^\top (x' - x_t)$, where $x' = \mathcal{P}_S(x_t - \beta^s g_t)$.
3. Set $x_{t+1} = \mathcal{P}_S(x_t - \beta^s g_t)$, $t \leftarrow t + 1$ and repeat steps 1 and 2, until the desired accuracy is achieved. Output $\hat{x} = x_t$.

Figure 2: The projected gradient descent algorithm.

The gradient $\nabla_\pi f(\pi; X)$ in Eq. (3) is easy to compute:

$$\frac{\partial f(\pi; X)}{\partial \pi_i} = -x_i^\top (X^\top \text{diag}(\pi) X)^{-1} x_i, \qquad i = 1, \cdots, n.$$

Because $X^\top \text{diag}(\pi) X$ is a shared term, computing $\nabla_\pi f(\pi; X)$ takes $O(np^2 + p^3)$ operations. The projection step onto the intersection of ℓ_1 and ℓ_∞ balls is complicated and non-trivial, which we describe in details in the supplementary material. In general, the projection step can be done in $O(n \log \|\pi\|_\infty)$ time, where π is the point to be projected.

The following proposition establishes convergence guarantee for the projected gradient descent algorithm. Its proof is given in the supplementary material.

Proposition B.1. *Let $\pi^{(0)}$ be the "flat" initialization (i.e., $\pi_i^{(0)} = k/n$) and $\pi^{(t)}$ be the solution after t projected gradient iterations. Then*

$$f(\pi^{(t)}; X) - f(\pi^*; X) \leq \frac{\|X\|_2^4 [\mathrm{tr}(\hat{\Sigma}^{-1})]^3 \|\pi^{(0)} - \pi^*\|_2^2}{\beta t},$$

where $\hat{\Sigma} = \frac{1}{n} X^\top X$.

We also remark that the provided convergence speed is very conservative, especially in the cases when k is large where practical evidence suggests that the algorithm converges in very few iterations (cf. Sec. 4).

C Fedorov's exchange algorithm

The algorithm starts with a initial subset $S \subseteq [n]$, $|S| \leq k$, usually initialized with random indices. A "best" pair of exchanging indices are computed as

$$i^*, j^* = \mathrm{argmin}_{i \in S, j \notin S} \ell(S \backslash \{i\} \cup \{j\}),$$

where $\ell(S)$ is the objective function to be minimized. In our case it would be the A-optimality objective $\ell(S) = F(S; X) = \mathrm{tr}((X_S^\top X_S)^{-1})$. The algorithm then "exchanges" i^* and j^* by setting $S' \leftarrow S \backslash \{i^*\} \cup \{j^*\}$ and continues such exchanges until no exchange can lower the objective value. Under without replacement settings, special care needs to be taken to ensure that S consists of distinct indices.

Computing the objective function $F(S; X) = \mathrm{tr}((X_S^\top X_S)^{-1})$ requires inverting a $p \times p$ matrix, which could be computationally slow. A more computationally efficient approach is to perform rank-1 update of the inverse of $X_S^\top X_S$ after each iteration, via the Sherman-Morrison formula:

$$(A + uv^\top)^{-1} = A^{-1} - \frac{A^{-1} u v^\top A^{-1}}{1 + u^\top A^{-1} v}.$$

Each exchange would then take $O(nkp^2)$ operations. The total number of exchanges, however, is unbounded and could be as large as $\binom{n}{k}$ in theory. In practice we do observe that a large number of exchanges are required in order to find a local optimal solution.

References

Ageev, A. A., & Sviridenko, M. I. (2004). Pipage rounding: A new method of constructing algorithms with proven performance guarantee. *J. Comb. Optim.*, *8*(3), 307–328.

Anderson, D. G., Gu, M., & Melgaard, C. (2014). An efficient algorithm for unweighted spectral graph sparsification. *arXiv preprint arXiv:1410.4273*.

Avron, H., & Boutsidis, C. (2013). Faster subset selection for matrices and applications. *SIAM J. Matrix Anal. Appl.*, *34*(4), 1464–1499.

Batson, J., Spielman, D. A., & Srivastava, N. (2012). Twice-ramanujan sparsifiers. *SIAM J. Comput.*, *41*(6), 1704–1721.

Chaudhuri, K., Kakade, S., Netrapalli, P., & Sanghavi, S. (2015). Convergence rates of active learning for maximum likelihood estimation. In *Proceedings of Advances in Neural Information Processing Systems*.

Chen, S., Varma, R., Singh, A., & Kovačević, J. (2015). Signal representations on graphs: Tools and applications. *arXiv:1512.05406*.

Chernoff, H. (1953). Locally optimal designs for estimating parameters. *Ann. Statist.*, (pp. 586–602).

Davenport, M. A., Massimino, A. K., Needell, D., & Woolf, T. (2015). Constrained adaptive sensing. *IEEE Trans. Signal Process*, *64*(20), 5437–5449.

Dhillon, P., Lu, Y., Foster, D., & Ungar, L. (2013). New sampling algorithms for fast least squares regression. In *Proceedings of Advances in Neural Information Processing Systems*.

Drineas, P., Mahoney, M. W., & Muthukrishnan, S. (2008). Relative-error CUR matrix decompositions. *SIAM J. Matrix Anal. Appl.*, *30*(2), 844–881.

Drineas, P., Mahoney, M. W., Muthukrishnan, S., & Sarlós, T. (2011). Faster least squares approximation. *Numer. Math.*, *117*(2), 219–249.

Ein-Dor, P., & Feldmesser, J. (1987). Attributes of the performance of central processing units: a relative performance prediction model. *Communications of the ACM*, *30*(4), 308–317.

Hazan, E., & Karnin, Z. (2015). Hard margin active linear regression. In *Proceedings of International Conference on Machine Learning*.

Horel, T., Ioannidis, S., & Muthukrishnan, S. (2014). Budget feasible mechanisms for experimental design. In *Latin American Symposium on Theoretical Informatics*, (pp. 719–730). Springer.

Khuri, A., Mukherjee, B., Sinha, B., & Ghosh, M. (2006). Design issues for generalized linear models: a review. *Statist. Sci.*, *21*(3), 376–399.

Leeds, D. D., Pyles, J. A., & Tarr, M. J. (2014). Exploration of complex visual feature spaces for object perception. *Frontiers in computational neuroscience*, *8*, 106.

Ma, P., Mahoney, M. W., & Yu, B. (2015). A statistical perspective on algorithmic leveraging. *J. Mach. Learn. Res.*, *16*(1), 861–911.

Marcus, A., Spielman, D., & Srivastava, N. (2015a). Interlacing families I: bipartite ramanujan graphs of all degrees. *Ann. of Math.*, *182*, 307–325.

Marcus, A., Spielman, D., & Srivastava, N. (2015b). Interlacing families II: Mixed characteristic polynomials and the kadison-singer problem. *Ann. of Math.*, *182*, 327–350.

Miller, A., & Nguyen, N.-K. (1994). A fedorov exchange algorithm for d-optimal design. *J. R. Stat. Soc. Ser. C. Appl. Stat.*, *43*(4), 669–677.

Pukelsheim, F. (1993). *Optimal design of experiments*, vol. 50. SIAM.

Raskutti, G., & Mahoney, M. (2015). Statistical and algorithmic perspectives on randomized sketching for ordinary least-squares. In *Proceedings of International Conference on Machine Learning*.

Rudelson, M., & Vershynin, R. (2007). Sampling from large matrices: An approach through geometric functional analysis. *J. ACM*, *54*(4), 21.

Sabato, S., & Munos, R. (2014). Active regression by stratification. In *Proceedings of Advances in Neural Information Processing Systems*.

Spielman, D. A., & Srivastava, N. (2011). Graph sparsification by effective resistances. *SIAM J. Comput.*, *40*(6), 1913–1926.

Tibshirani, R. (2013). The lasso problem and uniqueness. *Electron. J. Stat.*, *7*, 1456–1490.

Van der Vaart, A. W. (2000). *Asymptotic statistics*, vol. 3. Cambridge university press.

Vandenberghe, L., & Boyd, S. (1996). Semidefinite programming. *SIAM Rev.*, *38*(1), 49–95.

Woodruff, D. (2014). Sketching as a tool for numerical linear algebra. *Found. Trends Theor. Comput. Sci.*, *10*(1–2), 1–157.

Yu, A. W., Su, H., & Li, F. (2012). Efficient euclidean projections onto the intersection of norm balls. In *Proceedings of International Conference on Machine Learning*.

Zhu, R., Ma, P., Mahoney, M., & Yu, B. (2015). Optimal subsampling approaches for large sample linear regression. *arXiv:1509.05111*.

Supplementary Material for: On Computationally Tractable Selection of Experiments in Regression Models

Yining Wang[1], Adams Wei Yu[1], and Aarti Singh[1]

[1]Machine Learning Department, Carnegie Mellon University

December 5, 2016

In Sec. A We describe detail implementation of the projection step onto the intersectio of ℓ_1 and ℓ_∞ balls, which is used in the projected gradient descent algorithm for solving Eq. (3). We also give a convergence analysis in Sec. B, which shows that projected gradient descent computes the global optimal solution of Eq. (3) to arbitrary precision with polynomial running time.

A Projection onto the intersection of ℓ_1 and ℓ_∞ norm balls

Similar to Yu et al. (2012), we only need to consider the case that π lies in the first quadrant, i.e., $\pi_i \geq 0, i = 1, 2, ..., n$. Then the projection x also lies in the first quadrant. Furthermore, we assume the point to be project lies in the area which is out of both norm balls and projection purely onto either ℓ_1 or ℓ_∞ ball is not the intersection of both. Otherwise, it is trivial to conduct the projection. The projection problem is formulated as follows:

$$\begin{aligned} \min_{x} \quad & \frac{1}{2}\|\pi - x\|_2^2 \\ \text{s.t.} \quad & \|x\|_1 \leq c_1, \ \ \|x\|_\infty \leq c_2. \end{aligned}$$

By introducing an auxiliary variable d, the problem above has the following equivalent form:

$$\begin{aligned} \min_{x,d} \quad & \frac{1}{2}\|\pi - x\|_2^2 \\ \text{s.t.} \quad & \|x\|_1 \leq c_1, \\ & x_i \leq d, \ i = 1, 2, ...n, \\ & d \leq c_2. \end{aligned} \tag{S1}$$

The Lagrangian of problem (S1) is

$$\mathcal{L}(x, d, \lambda_1, \lambda_2) = \frac{1}{2}||\pi - x||_2^2 + \sum_i \omega_i(x_i - d) + \lambda_1(\|x\|_1 - c_1) + \lambda_2(d - c_2) \tag{S2}$$

input : π, c_1, c_2 and precision parameter δ.
output: x, the projection of π onto $\{x : \|x\|_1 \leq c_1\} \cap \{x : \|x\|_\infty \leq c_2\}$.
Initialization: $\lambda_1 = \|\pi\|_\infty$.
while $|h(\lambda_1)| > \delta$ **do**
  if $h(\lambda_1) > \delta$ **then**
    $l = \lambda_1$;
  else
    $r = \lambda_1$;
  end
  $\lambda_1 = (l + r)/2$;
  for each i **do** $x_i = \min(\max(\pi_i - \lambda_1, 0), c_2)$;
end

Algorithm 3: Projecting π onto $\{x : \|x\|_1 \leq c_1\} \cap \{x : \|x\|_\infty \leq c_2\}$.

Let x^*, d^* and λ_1^*, λ_2^* respectively be the primal and dual solution of (S1), then its KKT condition is:

$$0 \in \partial \mathcal{L}_x(x^*, d^*, \lambda_1^*, \lambda_2^*), \tag{S3}$$
$$0 \in \partial \mathcal{L}_d(x^*, d^*, \lambda_1^*, \lambda_2^*), \tag{S4}$$
$$w_i^*(x_i^* - d^*) = 0, \quad i = 1, 2, ..., n \tag{S5}$$
$$\lambda_2^*(d^* - c_2) = 0, \tag{S6}$$
$$\lambda_1^*(\|x^*\|_1 - c_1) = 0, \tag{S7}$$
$$x_i^* \leq d^*, \quad i = 1, 2, ..., n \tag{S8}$$
$$d^* \leq c_2, \tag{S9}$$
$$\|x^*\|_1 \leq c_1, \tag{S10}$$
$$x_i^*, d^*, \lambda_1^*, \lambda_2^*, w_i^* \geq 0. \tag{S11}$$

The following lemmas illustrate the relation between the primal and dual solution of problem (S1):

Lemma A.1. *Either one of the following holds: 1)* $d^* > 0$ *and* $\sum_i (\max(\pi_i - \lambda_1^*, 0) - x_i^*) = \lambda_2^*$*; 2)* $d^* = 0$ *and* $\sum_i \max(\pi_i - \lambda_1^*, 0) \leq \lambda_2^*$.

Lemma A.2. *For* $i = 1, 2, ..., n$*, the optimal* x_i^* *satisfies*

$$x_i^* = \min(\max(\pi_i - \lambda_1^*, 0), d^*). \tag{S12}$$

Proof. Lemma A.1 and Lemma A.2 can be both viewed as the special cases of Lemma 6 and 7 in Yu et al. (2012). □

Now the KKT conditions can be reduced to finding $\lambda_1^*, \lambda_2^*, d^*$ that satisfy the following equa-

tions:

$$\sum_i \min(\max(\pi_i - \lambda_1^*, 0), d^*) - c_1 = 0, \tag{S13}$$

$$d^* - c_2 = 0, \tag{S14}$$

$$\sum_i \max(\max(\pi_i - \lambda_1^*, 0) - d^*, 0) = \lambda_2^* \text{ iff } d^* > 0, \tag{S15}$$

$$\sum_i \max(\pi_i - \lambda_1^*, 0) \leq \lambda_2^* \text{ iff } d^* = 0, \tag{S16}$$

$$d^* \geq 0, \lambda_2^* \geq 0, \lambda_1^* \geq 0. \tag{S17}$$

Further simplification induces the following result:

Lemma A.3. *Suppose d^* and λ_1^* are the primal and dual solution respectively, then*

$$d^* = c_2, \tag{S18}$$

$$x_i^* = \min(\max(\pi_i - \lambda_1^*, 0), c_2). \tag{S19}$$

Proof. Direct result of (S12) and (S14). □

(S19) shows that the solution x^* is determined once the optimal λ_1^* is found. Given a $\lambda_1 \in [0, \max_i(\pi_i)]$, we only need to check whether (S13) holds by looking into the function

$$h(\lambda_1) = \sum_i \min(\max(\pi_i - \lambda_1, 0), c_2) - c_1, \tag{S20}$$

and λ_1^* is simply the zero point of $h(\cdot)$. The following theorem shows that $h(\cdot)$ is a strictly monotonically decreasing function, so a binary search is sufficient to find λ_1^*, and x_i^* can be determined accordingly.

Theorem A.1. *1) $h(\lambda_1)$ is a continuous piecewise linear function in $[0, \max_i\{\pi_i\}]$; 2) $h(\lambda_1)$ is strictly monotonically decreasing and it has a unique root in $[0, \max_i\{\pi_i\}]$.*

Proof. 1) is obviously true. It is easy to check that $h'(\lambda_1) < 0$ in each piece, $h(\max_i\{\pi_i\}) < 0$ and $h(0) > 0$, so 2) also holds. □

Complexity Analysis: Algorithm 3 is proposed based on the observation above to solve the projection problem, which is essentially a bisection to search λ_1. Given a precision δ, the iteration complexity of bisection is $O(\log(\max_i(\pi_i)/\delta))$. Besides, the time complexity of evaluating $h(\cdot)$ is $O(n)$, so the total complexity of Algorithm 3 is $O(n \log(\max_i(\pi_i)/\delta))$.

B Convergence analysis of PGD

We provide a convergence analysis of the projected gradient descent algorithm used in optimizing Eq. (3). The analysis shows that the PGD algorithm approximately computes the global optimum of Eq. (3) in polynomial time. In simulation studies, much fewer iterations are required for convergence than predicted by the theoretical results.

Because we're using exact projected gradient descent algorithms, the objective function shall decay monotonically and hence convergence of such algorithms can be established by showing Lipschitz continuity of ∇f on a specific *level set*; that is, for some Lipschitz constant $L > 0$ the following holds for all π, π' such that $f(\pi), f(\pi') \leq f(\pi^{(0)})$:

$$\|\nabla f(\pi) - \nabla f(\pi')\|_2 \leq L\|\pi - \pi'\|_2, \tag{S21}$$

where $\pi^{(0)} = (1/n, 1/n, \cdots, 1/n)$ is the initialization point. Once Eq. (S21) holds, linear convergence (i.e., $f(\pi^{(t)}) - f(\pi^*) = O(1/t)$) can be established via standard projected gradient analysis.

The main idea of establishing Eq. (S21) is to upper bound the spectral norm of the Hessian matrix $H = \nabla^2 f(\pi)$ *uniformly* over all points π that satisfies $f(\pi) \leq f(\pi^{(0)})$. As a first step, we derive analytic forms of H in the following proposition:

Proposition B.1. *Let $\widetilde{\Sigma} = X^\top \mathrm{diag}(\pi) X$. We then have that*

$$H = 2(X^\top \widetilde{\Sigma}^{-2} X) \circ (X^\top \widetilde{\Sigma}^{-1} X),$$

where $\circ$ denotes the element-wise Hadamard product between two matrices of same dimensions.

Proof. We first derive the gradient of f. Fix arbitrary $i \in [n]$. The partial derivative can be computed as

$$\frac{\partial f}{\partial \pi_i} = \frac{\partial \mathrm{tr}(\widetilde{\Sigma}^{-1})}{\partial \pi_i} = \left\langle \frac{\partial \mathrm{tr}(\widetilde{\Sigma}^{-1})}{\partial \widetilde{\Sigma}}, \frac{\partial \widetilde{\Sigma}}{\partial \pi_i} \right\rangle = \left\langle -\widetilde{\Sigma}^{-2}, x_i x_i^\top \right\rangle = -x_i^\top \widetilde{\Sigma}^{-2} x_i.$$

Here $\langle A, B \rangle = \mathrm{tr}(B^\top A)$ is the element-wise multiplication inner product between two matrices. The second-order partial derivatives can then be computed as

$$\begin{aligned}\frac{\partial^2 f}{\partial \pi_i \partial \pi_j} = -\frac{\partial (x_i^\top \widetilde{\Sigma}^{-2} x_i)}{\partial \pi_j} &= -x_i^\top \frac{\partial \widetilde{\Sigma}^{-2}}{\partial \pi_j} x_i = x_i^\top \widetilde{\Sigma}^{-2} \frac{\partial \widetilde{\Sigma}^2}{\partial \pi_j} \widetilde{\Sigma}^{-2} x_i \\ &= x_i^\top \widetilde{\Sigma}^{-2} (x_j x_j^\top \widetilde{\Sigma} + \widetilde{\Sigma} x_j x_j^\top) \widetilde{\Sigma}^{-2} x_i = 2(x_i^\top \widetilde{\Sigma}^{-2} x_j) \cdot (x_i^\top \widetilde{\Sigma}^{-1} x_j).\end{aligned}$$

Subsequently,

$$H = \nabla^2 f = \left[\frac{\partial^2 f}{\partial \pi_i \partial \pi_j} \right]_{i,j=1}^n = 2(X^\top \widetilde{\Sigma}^{-2} X) \circ (X^\top \widetilde{\Sigma}^{-1} X).$$

□

Corollary B.1. *Suppose π satisfies $f(\pi) \leq f(\pi^{(0)}) = \mathrm{tr}(\Sigma_0^{-1})$, where $\Sigma_0 = \frac{1}{n} X^\top X$. We then have*

$$\|H\|_2 = \|\nabla^2 f(\pi)\|_2 \leq 2\|X\|_2^4 \cdot [\mathrm{tr}(\Sigma_0^{-1})]^3.$$

Proof. Let $\rho(A) = \max_i |\sigma_i(A)|$ denote the *spectral range* of matrix A. Clearly, $\|A\|_2 \leq \rho(A)$ and $\|A\|_2 = \rho(A)$ for positive semi-definite matrices. In Hom & Johnson (1991) it is established that $\rho(A \circ B) \leq \rho(A)\rho(B)$. Subsequently,

$$\|H\|_2 \leq 2\|X\|_2^4 \|\widetilde{\Sigma}^{-1}\|_2^3 \leq 2\|X\|_2^4 \cdot [\mathrm{tr}(\widetilde{\Sigma}^{-1})]^3 \leq 2\|X\|_2^4 \cdot [\mathrm{tr}(\Sigma_0^{-1})]^3,$$

where the last inequality is due to the condition that $\mathrm{tr}(\widetilde{\Sigma}^{-1}) = f(\pi) \leq f(\pi^{(0)}) = \mathrm{tr}(\Sigma_0^{-1})$. □

From Corollary B.1, we can prove the convergence of the optimization procedures outlined in Appendix. B following standard analysis of projected gradient descent on objective functions with Lipschitz continuous gradient:

Theorem B.1. *Suppose $\pi^{(t)}$ is the solution at the tth iteration of the projected gradient descent algorithm and π^* is the optimal solution. We then have*

$$f(\pi^{(t)}) - f(\pi^*) \leq \frac{\|X\|_2^4 \cdot [\mathrm{tr}(\Sigma_0^{-1})]^3 \cdot \|\pi^{(0)} - \pi^*\|_2^2}{\beta t},$$

where $\beta \in (0, 1)$ is the backtracking parameter in backtracking line search.

References

Hom, R., & Johnson, C. (1991). *Topics in Matrix Analysis*. Cambridge UP, New York.

Yu, A. W., Su, H., & Li, F. (2012). Efficient euclidean projections onto the intersection of norm balls. In *Proceedings of International Conference on Machine Learning*.